Docket No.: J&R-1062

I hereby certify that this correspondence is being deposited with the United States Postal Service as First Class Mail in an envelope addressed to the Commissioner for Patents, P.O. Box 1450, Alexandria, VA 22313-1450 on the date indicated below.

By: _____ Date: July 17, 2003

IN THE UNITED STATES PATENT AND TRADEMARK OFFICE

Applicant : Gupta Abhay, et al.
Applic. No. : 10/600,554
Filed : June 20, 2003
Title : Configuration and Method Having a First Device and a Second Device Connected to the First Device Through a Cross Bar

CLAIM FOR PRIORITY

Commissioner for Patents,
P.O. Box 1450, Alexandria, VA 22313-1450

Sir:

Claim is hereby made for a right of priority under Title 35, U.S. Code, Section 119, based upon the European Patent Application 020 13 702.2, filed June 20, 2002.

A certified copy of the above-mentioned foreign patent application is being submitted herewith.

Respectfully submitted,

_____
For Applicant

LAURENCE A. GREENBERG
REG. NO. 29,308

Date: July 17, 2003

Lerner and Greenberg, P.A.
Post Office Box 2480
Hollywood, FL 33022-2480
Tel: (954) 925-1100
Fax: (954) 925-1101

/av

Europäisches Patentamt	European Patent Office	Office européen des brevets

Bescheinigung | Certificate | Attestation

Die angehefteten Unterlagen stimmen mit der ursprünglich eingereichten Fassung der auf dem nächsten Blatt bezeichneten europäischen Patentanmeldung überein.

The attached documents are exact copies of the European patent application described on the following page, as originally filed.

Les documents fixés à cette attestation sont conformes à la version initialement déposée de la demande de brevet européen spécifiée à la page suivante.

Patentanmeldung Nr. | Patent application No. | Demande de brevet n°

02013702.2

Der Präsident des Europäischen Patentamts;
Im Auftrag

For the President of the European Patent Office

Le Président de l'Office européen des brevets
p.o.

R C van Dijk

Anmeldung Nr:
Application no.: 02013702.2
Demande no:

Anmeldetag:
Date of filing: 20.06.02
Date de dépôt:

Anmelder/Applicant(s)/Demandeur(s):

Infineon Technologies AG
St.-Martin-Strasse 53
81669 München
ALLEMAGNE

Bezeichnung der Erfindung/Title of the invention/Titre de l'invention:
(Falls die Bezeichnung der Erfindung nicht angegeben ist, siehe Beschreibung.
If no title is shown please refer to the description.
Si aucun titre n'est indiqué se referer à la description.)

Arrangement having a first device, and a second device which is connected to the first device via a cross bar

In Anspruch genommene Priorität(en) / Priority(ies) claimed /Priorité(s) revendiquée(s)
Staat/Tag/Aktenzeichen/State/Date/File no./Pays/Date/Numéro de dépôt:

Internationale Patentklassifikation/International Patent Classification/ Classification internationale des brevets:

G06F13/00

Am Anmeldetag benannte Vertragstaaten/Contracting states designated at date of filing/Etats contractants désignées lors du dépôt:

AT BE CH CY DE DK ES FI FR GB GR IE IT LI LU MC NL PT SE TR

JANNIG & REPKOW
PATENTANWÄLTE
EUROPEAN PATENT AND TRADEMARK ATTORNEYS
DIPL.-ING. (UNIV.) PETER JANNIG
DR.-ING. DIPL.-ING. INES REPKOW
KLAUSENBERG 20, D-86199 AUGSBURG
DEUTSCHLAND · GERMANY
TEL. 08 21 / 9 81 93 TELEX 210 352 JUR D
FAX 08 21 / 9 81 95 EMAIL J-U-R@T-ONLINE.DE

PATENTANWÄLTE JANNIG & REPKOW · KLAUSENBERG 20 · D-86199 AUGSBURG

AUGSBURG, June 20, 2002

European Patent Application

Our Ref.: 0841 EP/P

Applicant: Infineon Technologies AG
St.-Martin-Str. 53
D-81669 München

ARRANGEMENT HAVING A FIRST DEVICE, AND A SECOND DEVICE WHICH IS CONNECTED TO THE FIRST DEVICE VIA A CROSS BAR

Description

Arrangement having a first device, and a second device which is connected to the first device via a cross bar

The present invention relates to an apparatus as claimed in the preamble of patent claim 1, that is to say to an arrangement having a first device, and having a second device which is connected to the first device via a cross bar.

An arrangement such as this may, but need not, be entirely or partially a component of a programmable unit such as a microprocessor, microcontroller, signal processor or the like.

The first device is a device which is referred to in the following text as a master unit and can initiate a transfer of data from or to the second device, that is to say a read or write access to the second device; it may therefore be, by way of example, but need not necessarily be, a CPU or a DMA controller of a programmable unit.

The second device is a device which is referred to as a slave unit in the following text and emits data requested by the master unit to the master unit, or receives data supplied to it from the master unit and further processes or stores this data; it may thus, by way of example, but need not necessarily be, a memory.

The master unit and the slave unit need not be connected directly to the cross bar. The connection can also be provided via a bus interface, a bus bridge or some other interface.

Normally, not only a master unit and a slave unit, but a number of master units and/or a number of slave units, are connected to a cross bar, and can be connected to one another via the cross bar.

The fundamental design of an arrangement such as this is shown in Figure 2.

The arrangement shown in Figure 2 comprises a first master unit M11, a second master unit M12, a third master unit M13, a first slave unit S11, a second slave unit S12, a third slave unit S13 and a cross bar XB1.

The master units M11 to M13 and the slave units S11 to S13 are connected to one another via the cross bar XB1. To be more precise, this is done in such a way that the master units M11 to M13 and the slave units S11 to S13 are connected by means of lines or buses, which are not shown in any more detail in Figure 2, to associated connections of the cross bar XB1, and such that the connections of the cross bar XB1 to which the master units M11 to M13 are connected are each connected to all the connections to which the slave units S11 to S13 are connected.

In addition to the internal connections which have been mentioned, the cross bar XB1 contains arbiters A11 to A13 and multiplexers MUX11 to MUX13.

The arbiters A11 to A13 are connected upstream of the connections of the cross bar XB1 to which the slave units S11 to S13 are connected. To be more precise, this is done in such a way

- that the arbiter A11 is connected upstream of that connection of the cross bar XB1 to which the slave unit S11 is connected,

- that the arbiter A12 is connected upstream of that connection of the cross bar XB1 to which the slave unit S12 is connected, and

- that the arbiter A13 is connected upstream of that connection of the cross bar XB1 to which the slave unit S13 is connected.

The arbiters A11 to A13 monitor whether any of the master units M11 to M13 are requesting a connection for the slave unit which is connected to that connection of the cross bar which is connected upstream of the respective arbiter, and produce a connection between the relevant slave unit and the master unit which has requested the connection, when an appropriate connection request is present and the slave unit is not currently connected to any other master unit or - for whatever reason - must be connected to another master unit prior to this.

The multiplexers MUX11 to MUX13 are connected upstream of those connections of the cross bar XB1 to which the master units M11 to M13 are connected. To be more precise, this is done such that:

- the multiplexer MUX11 is connected upstream of that connection of the cross bar XB1 to which the master unit M11 is connected,

- the multiplexer MUX12 is connected upstream of that connection of the cross bar XB1 to which the master unit M12 is connected, and

- the multiplexer MUX13 is connected upstream of that connection of the cross bar XB1 to which the master unit M13 is connected.

The multiplexers MUX11 to MUX13 are controlled by the arbiters A11 to A13, to be precise in such a manner that data which is emitted from the slave units is in each case supplied to the master unit, to be precise only that master unit

which has requested the connection for the relevant slave unit.

For the sake of completeness, it should be mentioned that at least those lines by means of which the master units M11 to M13 request a connection for one of the slave units S11 to S13 are not routed via the multiplexers.

It is also possible for different master units to be connected to different slave units at the same time. For example, the first master unit M11 can be connected to the second slave unit S12, the second master unit M12 can be connected to the first slave unit S11, and the third master unit M13 can be connected to the third slave unit S13 at the same time via those internal connections of the cross bar XB1 which are shown by thicker lines.

The cross bar XB1 thus allows data to be transmitted very efficiently between the devices connected to it.

However, this is true only when the execution of mutually corresponding actions, which one master unit can request from different slave units, take place from the point of view of the master unit in accordance with the same scheme, in particular having the same timing.

For example, this is not the case when the master unit receives the data requested from a first slave unit after n clock cycles and receives the data requested from a second slave unit later, that is to say only after n + m clock cycles. This may occur, for example, when the second slave unit requires a longer time to emit the data requested from it than the first slave unit.
If differences such as these are present,

- the special features of the respective slave units, in particular the reaction times of the slave units, must

be set in the master unit to the requirements emitted from the master unit, or

- the cross bar must contain so-called wait state generators, which produce so-called wait states to signal to the master units that the slave unit has not yet reacted to the request from the master unit.

However, this makes the design and operation of the master units and of the cross bar more complex and complicated.

Furthermore, the various reaction times of the slave units to a request from a master unit do not depend only on the design of the slave unit but also on the signal delay times between the master units and the slave units.

The length of the signal delay times depends, inter alia, on the length of the connecting lines between the master units and/or the slave units and the cross bar, so that the signal delay times may differ considerably from one another.

Furthermore, poor signal delay times may make it necessary to insert one or more pipeline stages, in the form of flipflops for example, in the signal paths between certain master units and the cross bar and/or between certain slave units and the cross bar, and these pipeline stages may result in additional delays in the reaction of the slave units to a request from a master unit.

Additional delays may, furthermore, also occur as a result of the master units and the slave units not being connected directly to the cross bar, via bus interfaces, bus bridges or the like.

If such additional delays are present, and these delays are also intended to be taken into account by appropriate settings of the master units or of the wait state generators for

the cross bar, the design and operation of the master units and of the cross bar becomes even more complex and complicated.

Another solution to the problems caused by the additional delays is for the clock frequency at which the data is transmitted between the devices connected to the cross bar to be reduced sufficiently that the different signal delay times have no effect on the reaction times, and no pipeline stages are required either. However, in this case, the system operates more slowly than the speed at which it could actually operate.

The present invention is thus based on the object of finding a possible way in which the components of an arrangement of the type described above can cooperate efficiently, and can be combined in a flexible manner, with these components having a simple design and being simple to operate.

According to the invention, this object is achieved by the arrangement claimed in patent claim 1.

The arrangement according to the invention is distinguished

- when a read access to the second device occurs, the first device reads the data emitted from the second device, when it receives a ready signal which is produced by the second device and is supplied to the first device via the cross bar, and
- when a write access occurs from the first device to the second device,
  - the first device emits the data to be written to the second device, when it receives a ready signal which is produced by the second device and is supplied to the first device via the cross bar, and
  - the second device reads the data emitted from the first device, when it receives a data valid signal

which is produced by the first device and is supplied to the second device via the cross bar.

In an arrangement such as this, the slave unit signals to the master unit, and the master unit signals to the slave unit, that the action in each case expected from the respective unit has been carried out, so that the special precautions mentioned initially do not need to be carried out either in the master unit, in the cross bar or in the slave unit, which makes it possible for the master unit or the slave unit to carry out the actions which must be carried out, or to prevent those actions from being carried out, once a specific state has occurred in the slave unit or in the master unit.

The claimed arrangement thus allows the components of this arrangement to cooperate efficiently and to be combined flexibly, and said components having a simple design and being simple to operate.

Since most units which can be used as a slave unit intrinsically produce a ready signal or a signal which can be used as a ready signal, or can produce such a signal with little effort, and most units which can be used as a master unit intrinsically produce a data valid signal or a signal which can be used as a data valid signal, or can produce such a signal with little effort, the claimed arrangement can, furthermore, even be produced more simply and can be operated more simply than conventional arrangements of the type under discussion.

Advantageous developments of the invention can be found in the dependent claims, in the following description and in the figures.

The invention will be described in more detail in the following text using an exemplary embodiment and with reference to the figures, in which:

Figure 1 shows the design of the arrangement described in the following text, and

Figure 2 shows the design of the conventional arrangement described initially.

The arrangement described in the following text is in principle designed in the same way as the arrangement illustrated in Figure 2 and described initially with reference to this figure. This means that it has at least one master unit, at least one slave unit and a cross bar connecting these units to one another. Figure 1, with reference to which the arrangement introduced here will be described, shows only one master unit and two slave units, in contrast to Figure 2, for the sake of clarity.

Each of the available master units can access (via the cross bar) at least one of the slave units. In the example under consideration, the accesses are read accesses, by means of which a master unit reads from a slave unit data stored in that slave unit, and write accesses by means of which a master unit transmits to the slave unit data to be stored or to be further processed in that slave unit. The accesses which the master units make to the slave units may, however, also be any other desired types of access.

The described arrangement in the example under consideration is a component of a programmable unit such as a microprocessor, microcontroller or signal processor. However, there is no restriction to this. In particular, certain master units and/or slave units may at least partially also be provided outside the programmable unit and, furthermore, the described arrangement may also entirely or partially be a component of an integrated or non integrated circuit.

As in the case of the conventional arrangement described initially, it is also true in the case of the arrangement de-

scribed in the following text that the master units and the slave units need not be connected directly to the cross bar, but may also be connected to the cross bar via a bus interface, a bus bridge or the like.

The arrangement described in the following text has a number of differences in comparison to the arrangement described initially with reference to Figure 2, and these will now be described with reference to Figure 1.

The arrangement shown in Figure 1 contains a master unit M1, a first slave unit S1, a second slave unit S2, and a cross bar XB.

The cross bar XB contains arbiters A1 and A2, multiplexers MUX1, MUX2 and MUX3, as well as pipeline stages PS1 to PS7, with the pipeline stages PS1 to PS7 in the example under consideration being formed by registers, to be more precise by flipflops, although they could also be implemented in any other desired way. The arrangement and the operation of the components of the cross bar XB will be described in more detail later on.

The master unit M1 and the slave units S1 and S2 are connected to one another via the cross bar XB. To be more precise, this is done in such a way that:

- the master unit M1
  - is connected via a first address bus ADDR1 to an input connection of the arbiter A1 and to an input connection of the arbiter A2,
  - is connected via a first write data bus WRITE1 to an input connection of the multiplexer MUX2 and to an input connection of the multiplexer MUX3, and
  - is connected via a first read data bus READ1 to the output connection of the multiplexer MUX1,
- the slave unit S1

- is connected via a second address bus ADDR2 to the output connection of the arbiter A1,
- is connected via a second write data bus WRITE2 to the output connection of the multiplexer MUX2, and
- is connected via a second read data bus READ2 to an input connection of the multiplexer MUX1, and

- the slave unit S2
- is connected via a third address bus ADDR3 to the output connection of the arbiter A2,
- is connected via a third write data bus WRITE3 to the output connection of the multiplexer MUX3, and
- is connected via a third read data bus READ3 to an input connection of the multiplexer MUX1.

A further master unit would
- be connected via a fourth address bus to a further input connection of the arbiter A1 and to a further input connection of the arbiter A2,
- be connected via a fourth write data bus to a further input connection of the multiplexer MUX2 and to a further input connection of the multiplexer MUX3, and
- be connected via a fourth read data bus to the output connection of a further multiplexer, whose input connections are connected to the read data buses READ1 and READ2.

Expressed in general form, this is done in such a way that:

- a multiplexer associated with the relevant master unit is provided for each of the master units connected to the cross bar, and

- an arbiter associated with the relevant slave unit and a multiplexer associated with the relevant slave unit are provided for each of the slave units connected to the cross bar,

with

- the multiplexer associated with each master unit
  - connecting the input connections via read data buses to all the slave units that are present, and
  - connecting the output connection via a read data bus to the master unit,

- the multiplexer associated with each slave unit
  - connecting the input connections via write data buses to all the master units that are present, and
  - connecting the output connection via a write data bus to the slave unit, and
- the arbiter associated with each slave unit
  - connecting the input connections via address buses to all the master units which are present, and
  - connecting the output connection via an address bus to the slave unit.

The multiplexers that are present are controlled by the arbiters.

Addresses which are emitted from the master units as well as various control signals which will be described in more detail later on are transmitted via the address buses.

Data emitted from the slave units as well as various control signals which will likewise be described in more detail later on are transmitted via the read data buses.

Data emitted from the master units as well as various control signals which will likewise be described in more detail later on are transmitted via the write data buses.

The master units which are connected to the cross bar XB can access the slave units for reading and/or writing. A read

access to a slave unit allows the master unit to read data stored in that slave unit; a write access to a slave unit allows the master unit to transmit data to be stored or further processed in the slave unit to that slave unit.

The procedures which take place during a read access and the procedures which take place during a write access will be described in more detail in the following text with reference to Figure 1.

In the case of the read access which is described first of all, it is assumed that the master unit M1 wishes to read data from the slave unit S1.

The read access thus starts with the master unit M1 emitting an address, a read signal and a request signal via the address bus ADDR1, with the master unit M1

- using the request signal to signal that it wishes to access one of the slave units,
- using the read signal to signal that it wishes to read data from the relevant slave unit, and
- using the address to indicate the slave unit or the point within the relevant slave unit from which data should be read.

The address, the read signal and the request signal are supplied via the address bus ADDR1 both to the arbiter A1 and to the arbiter A2. Each of the arbiters A1 and A2 uses the request signal to identify that a master unit is requesting an access, and then uses the address to check whether the access is intended to be made to the slave unit which is associated with the respective arbiter. In this case, the arbiter A1 finds that the requested access is intended to be made to the slave unit S1 which is associated with that arbiter A1; the

arbiter A2 finds that no access is intended to be made to the slave unit S2 which is associated with that arbiter A2.

Once the arbiter A1 has found that a master unit wishes to access the slave unit S1, it first of all emits a grant signal via the address bus to that master unit which has requested access to that slave unit S1; the master unit which has requested access can be determined on the basis of the address bus via which the request signal requesting access was transmitted. Thus, in the example under consideration, the arbiter A1 transmits a grant signal via the address bus ADDR1 to the master unit M1.

The master unit M1 uses the grant signal supplied to it to identify that its access request has been accepted, and now waits for the transmission of the data requested from the slave unit S1. The master unit has no information as to when this data will be supplied, and the master unit is also not signaled by means of wait state cycles or the like that the data requested from the slave unit S1 is not yet available. Instead of this, the master unit waits until it is supplied with a ready signal, which will be described in more detail later on.

After emitting the grant signal to the master unit M1, or even at the same time as this, the arbiter A1 checks whether the slave unit S1 can be accessed at that time. This is the situation when the slave unit is not already being accessed at that time and when there are no access requests to be processed in advance. Access requests which are to be processed in advance may, by way of example, be access requests which arrived at the arbiter A1 earlier, or which have a higher allocated priority than the present access request.

If the arbiter A1 finds that the access requested by the master unit M1 to the slave unit S1 can be carried out, it transmits the read signal and that part of the address which

is required for addressing that slave unit S1 via the address bus ADDR2 to the slave unit S1.

Essentially at the same time, the arbiter A1 actuates the multiplexer MUX1 such that it passes on the data supplied to it from the slave unit S1 via the read data bus READ2 to the master unit M1, via the read data bus READ1. In addition, it is possible to provide for the arbiter A1 to actuate the multiplexer MUX2 such that the latter passes on the data supplied to it from the master unit M1 via the write data bus WRITE1 to the slave unit S1 via the write data bus WRITE2.

The slave unit S1 uses the data supplied to it via the address bus ADDR2 to identify that it should read and emit the data stored at the address supplied to. It then reads the data to be read, and emits this data together with a ready signal via the read data bus READ2. The data transmitted via the read data bus READ1 is passed on via the multiplexer MUX1 and the read data bus READ1 to the master unit M1.

As has already been explained above, the master unit M1 waits to receive the ready signal. The ready signal signals to the master unit M1 that the data already requested by the master unit M1 from the slave unit S1 is available and can now be read. The master unit M1 reads the data supplied to it via the read data bus READ1, thus ending the read access process.

In the example under consideration, the master unit M1 had to wait to receive the ready signal, that is to say, in the meantime, it did not carry out any further access to one of the slave units. However, the master unit can make further accesses to the slave units even before reception of the ready signal. In particular, it is feasible without any problems for the master unit M1 to request further accesses to the slave unit which it is currently accessing, even before receiving the ready signal. This means that, even before the termination of an access to the slave unit S1, the master

unit M1 can make one or more further accesses to that slave unit S1.

Since the master unit M1 is caused to receive a ready signal, which is produced by the slave unit and is passed on through the cross bar XB to the master unit, in order to read the data emitted from the slave unit, there is no need

- for the master unit to be informed of this information that has been set, defining the time at which the data emitted from the slave unit S1 should be read, or

- for a wait state generator to be provided in the cross bar, or anywhere else, which signals to the master unit by producing wait state cycles that the data requested from the slave unit S1 is not yet available for reading.

This in turn makes it possible for any desired number of pipeline stages to be inserted at any desired points, completely independently of one another, in the buses via which the units connected to the cross bar are connected to the cross bar, without any need for this to be taken into account in the design and configuration of the master unit M1 and of the cross bar XB. Figure 1 shows one possible arrangement of pipeline stages. As has already been mentioned above, the pipeline stages are annotated by the reference symbols PS1 to PS7. The pipeline stages may also be provided outside of the cross bar.

The pipeline stages result in the data and signals which are transmitted via the data and signal paths which contain the pipeline stages being transmitted delayed by one or more clock signal periods. However, in the case of the arrangement which is illustrated in Figure 1 and is described with reference to that figure, this cannot interfere with correct operation of the arrangement, in any circumstances. The master unit M1 reads the data supplied to it from the slave unit S1

when it receives the ready signal that is transmitted at the same time as the data, so that it is irrelevant how long it takes between requesting access to the slave unit S1 or receiving the grant signal and the reception of the data requested from the slave unit.

A corresponding situation also arises, of course, when the master unit M1 wishes to read data from another of the slave units that are present, or when another of the master units which are present wishes to read data from one slave unit.

A similar situation occurs when the master unit M1 makes a write access to the slave unit S1, that is to say when the master unit M1 transmits to the slave unit S1 data to be stored in it or to be further processed in it.

The write access starts with the master unit M1 emitting via the address bus ADDR1 an address, a write signal and a request signal, with the master unit M1

- signaling by means of the request signal that it wishes to access one of the slave units,

- signaling by means of the write signal that it wishes to write data to the relevant slave unit, and

- indicating by the address the slave unit or the point within the relevant slave unit to which the data should be written.

The address, the write signal and the request signal are supplied via the address bus ADDR1 both to the arbiter A1 and to the arbiter A2. Each of the arbiters A1 and A2 uses the request signal to identify that a master unit has requested access, and then uses the address to check whether the access is intended to be made to the slave unit which is associated with the respective arbiter. In this case, the arbiter A1

finds that the requested access is intended to be made to the slave unit S1 which is associated with the arbiter A1; the arbiter A2 finds that no access is intended to be made to the slave unit S2 which is associated with that arbiter A2.

Once the arbiter A1 has found that a master unit wishes to access the slave unit S1, it first of all emits, via the address bus, a grant signal to that master unit which has requested access to the slave unit S1; the master unit which has requested access can be determined on the basis of the address bus via which the request signal requesting access was transmitted. Thus, in the example under consideration, the arbiter A1 transmits a grant signal via the address bus ADDR1 to the master unit M1.

The master unit M1 uses the grant signal supplied to it to identify that its access request has been received, and now waits until it can emit the data to the slave unit S1 which it wishes to store in that slave unit S1. The master unit has no information about the time at which this can be done, and the master unit is also not signaled by means of wait state cycles or the like that the relevant time has not yet been reached. Instead of this, the master unit waits until it is supplied with a ready signal, which will be described in more detail later on.

After emitting the grant signal to the master unit M1, or even at the same time as this, the arbiter A1 checks whether access is currently possible to the slave unit S1. This is the case where no access is currently being made to the slave unit and when there are no access requests which need to be dealt with in advance. Access requests which need to be dealt with in advance may, for example, be access requests which arrived at the arbiter A1 earlier, or which have a higher associated priority than the current access request.

If the arbiter A1 finds that the access requested by the master unit M1 to the slave unit S1 can be carried out, it transmits the write signal and that part of the address which is required for addressing the slave unit S1, via the address bus ADDR2 to the slave unit S1.

Essentially at the same time:

- the arbiter A1 actuates the multiplexer MUX1 such that it passes on the data supplied to it from the slave unit S1 via the read data bus READ2, via the read data bus READ1, to the master unit M1, and

- the arbiter A1 actuates the multiplexer MUX2 such that it passes on the data supplied to it from the master unit M1 via the write data bus WRITE1, via the write data bus WRITE2, to the slave unit S1.

The slave unit S1 uses the data supplied to it via the address bus ADDR2 to identify that data should be stored at the address supplied to it. It then emits a ready signal on the read data bus READ2. This signal is passed via the multiplexer MUX1 and the read data bus READ2 to the master unit M1.

As has already been explained above, the master unit M1 waits to receive the ready signal. The ready signal signals to the master unit M1 that the master unit M1 may now emit the data to be written to the slave unit S1. The master unit M1 then emits the data to be written to the slave unit S1, together with a data valid signal, via the write data bus WRITE1. This data is passed via the multiplexer MUX2 and the write data bus WRITE2 to the slave unit S1. The slave unit S1 uses the data valid signal to identify that the data to be written to it is available. It reads this data, and stores it or processes it further. This completes the write access by the master unit M1 to the slave unit S1.

In the example under consideration, the master unit M1 had to wait to receive the ready signal, that is to say in the meantime it did not make any further access to one of the slave units. However, even before receiving the ready signal, the master unit can make further accesses to the slave units. In particular, it is possible without any problems for the master unit M1 to request further accesses to the slave unit which it is currently accessing, even before receiving the ready signal. This means that, even before completion of an access to the slave unit S1, the master unit M1 may request one or more further accesses to the slave unit S1.

Since the master unit M1 is caused to receive a ready signal, which is produced by the slave unit and is passed on through the cross bar XB to the master unit, in order to emit the data to be written to the slave unit S1, there is no need

- for the master unit to be informed of this information that has been set, defining the time at which the data emitted from the slave unit S1 should be read, or

- for a wait state generator to be provided in the cross bar, or anywhere else, which signals to the master unit by producing wait state cycles that the data to be written to the slave unit S1 is not yet to be emitted.

This in turn makes it possible for any desired number of pipeline stages to be inserted at any desired points, completely independently of one another, in the buses via which the units connected to the cross bar are connected to the cross bar, without any need for this to be taken into account in the design and configuration of the master unit M1 and of the cross bar XB. Figure 1 shows one possible arrangement of pipeline stages. As has already been mentioned above, the pipeline stages are annotated by the reference symbols PS1 to

PS7. The pipeline stages may also be provided outside of the cross bar.

The pipeline stages result in the data and signals which are transmitted via the data and signal paths which contain the pipeline stages being transmitted delayed by one or more clock signal periods. However, in the case of the arrangement which is illustrated in Figure 1 and is described with reference to that figure, this cannot interfere with correct operation of the arrangement, in any circumstances. On reception of the ready signal, the master unit M1 emits the data to be written to the slave unit S1, so that it is irrelevant how long it takes between the request for access to the slave unit S1 or the reception of the grant signal and the emission of the data to be written to the slave unit.

A corresponding situation also arises, of course, when the master unit M1 wishes to write data to another of the slave units which are present, or when another of the master units which are present wishes to write data to a slave unit.

The described arrangement can be implemented and operated in a simple manner, and allows the devices which are connected to one another via the cross bar to cooperate efficiently, and to be combined in a flexible manner.

Claims

1. An arrangement having a first device (M1), and a second device (S1) which is connected to the first device via a cross bar (XB), in which case the first device can access the second device via the cross bar for reading and/or writing, wherein

- when a read access to the second device occurs, the first device reads the data emitted from the second device, when it receives a ready signal which is produced by the second device and is supplied to the first device via the cross bar, and
- when a write access occurs from the first device to the second device,
  - the first device emits the data to be written to the second device, when it receives a ready signal which is produced by the second device and is supplied to the first device via the cross bar, and
  - the second device reads the data emitted from the first device, when it receives a data valid signal which is produced by the first device and is supplied to the second device via the cross bar.

2. The arrangement as claimed in claim 1, wherein the first device (M1) and the cross bar (XB) are connected to one another via a first address bus (ADDR1), a first read data bus (READ1), and a first write data bus (WRITE1), and wherein the second device (S1) and the cross bar (XB) are connected to one another via a second address bus (ADDR2), a second read data bus (READ2), and a second write data bus (WRITE2).

3. The arrangement as claimed in claim 1 or 2, wherein

the first device (M1) sends a request signal to the cross bar (XB) when it wishes to make a read access to the second device (S1).

4. The arrangement as claimed in claim 3, wherein the first device (M1) sends an address to the cross bar (XB) at the same time as the request signal, which address specifies the device, and the point within the device, from which data should be read.

5. The arrangement as claimed in claims 2 to 4, wherein the request signal and the address are transmitted to the cross bar (XB) via the first address bus (ADDR).

6. The arrangement as claimed in claim 3, wherein the cross bar (XB) confirms the read access request by transmission of a grant signal to the first device (M1).

7. The arrangement as claimed in claims 2 and 6, wherein the grant signal is transmitted to the first device (M1) via the first address bus (ADDR1).

8. The arrangement as claimed in claim 4, wherein the cross bar (XB) passes on at least a portion of the address supplied to it via the second address bus (ADDR2) to the device from which data should be read.

9. The arrangement as claimed in claim 8, wherein the second device (S1) emits to the cross bar (XB) the data which is stored at the address supplied to it.

10. The arrangement as claimed in claim 9,

wherein
the second device (S1) emits the ready signal to the cross bar (XB) at the same time that it emits the data that has been read.

11. The arrangement as claimed in claim 10,
wherein
the data that has been read and the ready signal are transmitted to the cross bar (XB) via the second read data bus (READ2).

12. The arrangement as claimed in claim 11,
wherein
the cross bar (XB) passes on the data supplied to it and the ready signal via the first read data bus (READ1) to the first device (M1).

13. The arrangement as claimed in one of the preceding claims,
wherein
the first device (M1) sends a request signal to the cross bar (XB) when it wishes to make a write access to the second device (S1).

14. The arrangement as claimed in claim 13,
wherein
the first device (M1) sends an address to the cross bar (XB) at the same time as the request signal, which address specifies the device and the point within this device to which data should be written.

15. The arrangement as claimed in claims 2, 13 and 14,
wherein
the request signal and the address are transmitted to the cross bar (XB) via the first address bus (ADDR1).

16. The arrangement as claimed in claim 13,

wherein
the cross bar (XB) confirms the write access request by transmission of a grant signal to the first device (M1).

17. The arrangement as claimed in claims 2 and 16, wherein
the grant signal is transmitted to the first device (M1) via the first address bus (ADDR1).

18. The arrangement as claimed in claim 14, wherein
the cross bar (XB) passes on at least a portion of the address supplied to it via the second address bus (ADDR1) to the device to which data should be written.

19. The arrangement as claimed in claim 18, wherein
the second device (S1) emits the ready signal to the cross bar (XB) when it is ready to receive the data to be stored in it.

20. The arrangement as claimed in claim 19, wherein
the ready signal is transmitted to the cross bar (XB) via the second read data bus (READ2).

21. The arrangement as claimed in claim 20, wherein
the cross bar (XB) passes on the ready signal to the first device (M1) via the first read data bus (READ1).

22. The arrangement as claimed in claim 21, wherein
the first device (M1) emits to the cross bar (XB) the data to be written to the second device (S1).

23. The arrangement as claimed in claim 22,

wherein
the first device (M1) emits the data valid signal to the cross bar (XB) at the same time that it emits the data to be written to the second device (S1).

24. The arrangement as claimed in claim 23,
wherein
the data emitted from the first device (M1) and the data valid signal are transmitted to the cross bar (XB) via a first write data bus (WRITE1).

25. The arrangement as claimed in claim 24,
wherein
the cross bar (XB) passes on the data supplied to it and the data valid signal to the second device (S1) via a second write data bus (WRITE2).

Abstract

Arrangement having a first device, and a second device which is connected to the first device via a cross bar

An arrangement is described, having a first device and a second device which is connected to the first device via a cross bar, in which case the first device can access the second device via the cross bar for reading and/or writing. The described arrangement is distinguished in that:

- when a read access to the second device occurs, the first device reads the data emitted from the second device, when it receives a ready signal which is produced by the second device and is supplied to the first device via the cross bar, and
- when a write access occurs from the first device to the second device,
  - the first device emits the data to be written to the second device, when it receives a ready signal which is produced by the second device and is supplied to the first device via the cross bar, and
  - the second device reads the data emitted from the first device, when it receives a data valid signal which is produced by the first device and is supplied to the second device via the cross bar.

Figure 1

List of reference symbols

Ax	Arbiter
ADDRx	Address bus
Mx	Master unit
MUXx	Multiplexer
PSx	Pipeline stage
READx	Read data bus
Sx	Slave unit
WRITEx	Write data bus
XB	Cross bar

200002767

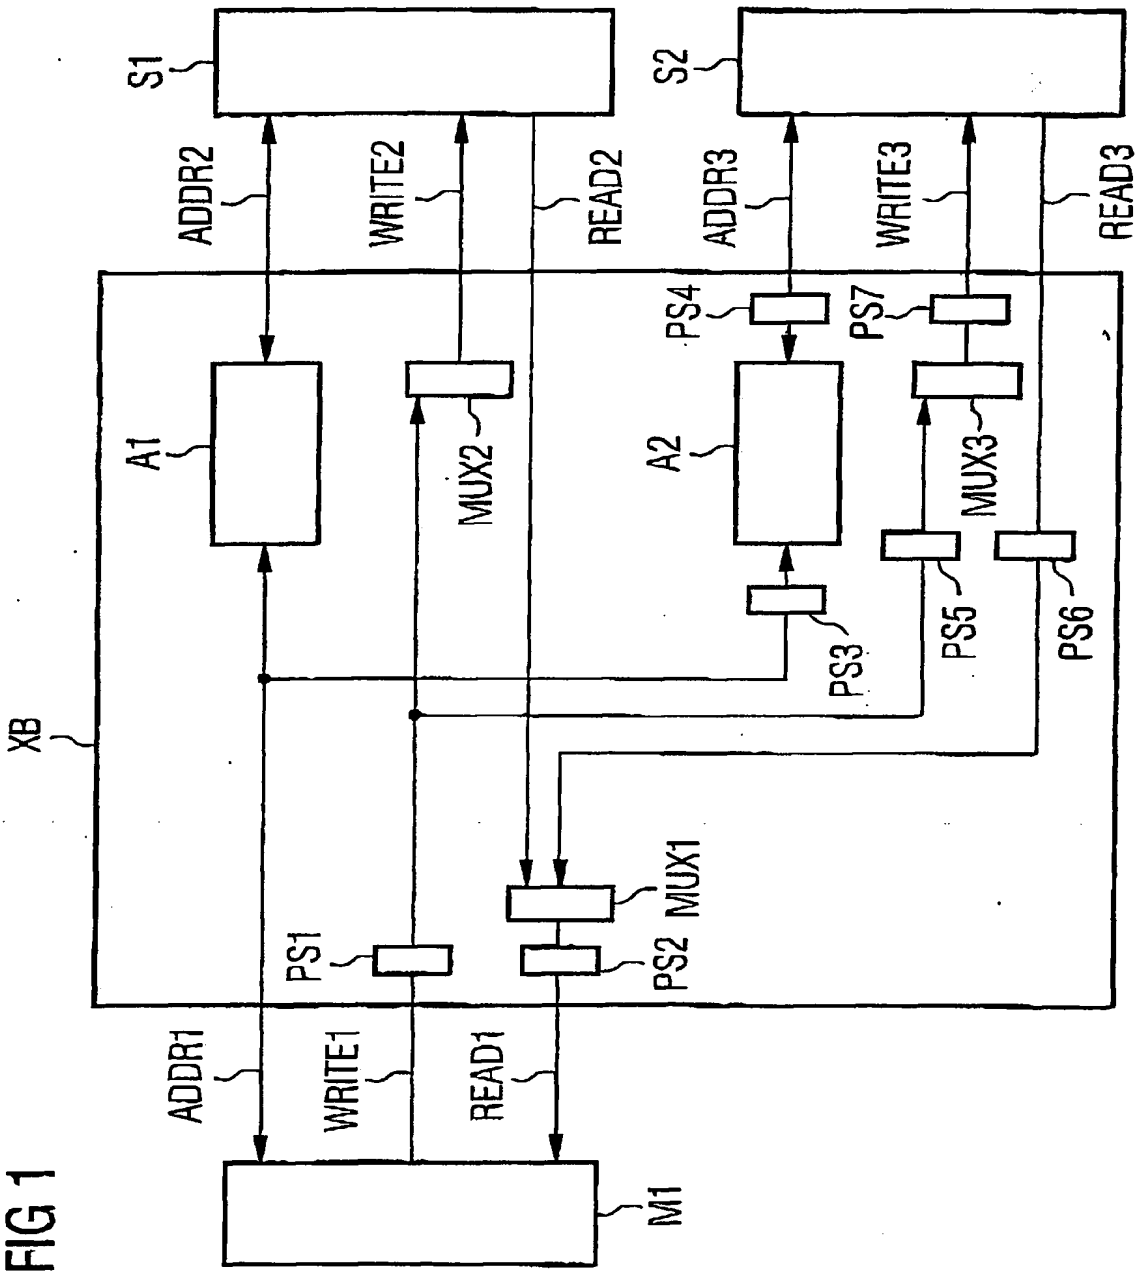

FIG 1

200002767

FIG 2